

# SINGULARITY-FREE BUNDLE ADJUSTMENT

E. Garcia

Airbus DS Geo SA, Sophia-Antipolis, France - emmanuel.garcia@airbus.com

**KEY WORDS:** Bundle adjustment, Structure-from-Motion

## ABSTRACT:

The photogrammetric bundle adjustment is well-behaved in the case of structured aerial imagery looking in the nadir direction. That is less so in the case of ground-level imagery with less structure and potentially looking in any direction. Besides, the cost function based on reprojection errors of tie points is not defined everywhere and exhibits singularities which renders this bundle adjustment process sensitive to initial conditions and outliers. In order to handle difficult configurations without incurring the risks posed by the reprojection function, we propose a new error function that is equivalent to the reprojection error when this error tends to zero, and that enjoys many desirable properties, such as being defined everywhere and being continuous. This allows an easier implementation of a robust bundle adjustment, and incidentally it also allows to solve derivative problems such as triangulating points starting from arbitrary initial positions, or estimating the relative positions of calibrated and oriented cameras starting from arbitrary positions, thus offering a simple solution to the known-orientation structure-from-motion problem.

## 1. INTRODUCTION

In computer vision a prerequisite to producing a 3D reconstruction is to accurately model the geometry of the images. This often involves an approximate initial knowledge of this geometry, followed by a global optimization called bundle adjustment (for an all-around review we refer to (Triggs et al., 2000)).

This process typically uses point correspondences between the images, and aims at minimizing the quadratic sum of reprojection errors of the triangulated points. As it turns out, this optimization problem may have many local minima, even in trivial configurations (Kahl and Hartley, 2008), and so is sensitive to initial conditions. Also, when the triangulated points are used as auxiliary variables, their projection in some images may not be defined, and/or singularities may arise in the course of the process. This may be countered by preemptively discarding some points, or discarding them during the optimization, using some *ad hoc* criteria.

Cost functions other than the quadratic sum of reprojection errors have been proposed in the literature for their practical advantage in optimization (Schinstock et al., 2009; Mitra and Chellappa, 2008), but then somewhat changing the meaning and significance of the optimum.

We propose a new cost function that has many desirable mathematical properties ensuring the robustness of the bundle adjustment process. Close to the optimum, this cost function is mathematically equivalent to the standard quadratic sum of reprojection errors, but contrary to the latter, it is well defined everywhere. Using this cost function the bundle adjustment problem can be implemented easily with not special cases to be taken care of. Also, it turns out that other optimization problems can be solved using this same bundle adjustment process, such as triangulating a point starting from an arbitrary location, or estimating relative camera positions starting from arbitrary positions.

The new cost function has been designed incrementally by adding mathematical requirements. In section 2 we describe this cost function and the rationale leading to it.

In section 3, we comment on bundle adjustment experiments using this cost function, and present results of the known-orientation structure-from-motion (SfM) problem using a simple bundle-adjustment-like optimization, being rendered possible by the new cost function.

## 2. DESIGN OF A COST FUNCTION

### 2.1 Shortcoming of the reprojection error

Given a 3D point  $M$  and its known position  $m$  in a given camera, the reprojection error is defined as  $r = P(M) - m$  where  $P : \mathbb{R}^3 \rightarrow I \subset \mathbb{R}^2$  is the projection function from object space to image space  $I$ . The cost function for this point is often defined as the squared norm of the reprojection error vector, the implicit assumption being then that the measurement error follows a Gaussian probability distribution. This allows using non-linear least squares optimization algorithms, such as Gauss-Newton or derivatives.

As not all points project inside the image, the projection function is not defined over the whole object space, which is a concern if this condition is not properly checked and handled during optimization, and doing so adds some complexity to a would-be robust bundle adjustment algorithm.

The main goal is thus to find an extension of the projection function to the whole object space that has good properties, so as to simplify a bundle adjustment implementation. If that is impossible, we want to find an alternative cost function that is defined over the whole object space, and that is equivalent to the squared norm of the reprojection error when it tends towards zero.

### 2.2 Tentative extension of the reprojection error

Assuming a simple pinhole model with no optical distortion, the formula for the reprojection error can be extended as is from the camera frustum to the entire open half-space in front of the camera. However it cannot be further extended continuously at the optical center of the camera, nor on the plane perpendicular

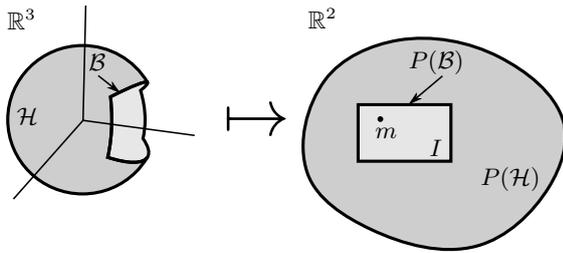


Figure 1. The projection of the hidden part  $\mathcal{H}$  of the unit sphere by a continuous function would necessarily contain any point  $m$  of the image domain  $I$ .

to the optical axis going through the optical center, and consequently not behind the camera either. For such an extension to be continuous over the whole space, it must of necessity be bounded on any bounded region of space (in particular a sphere centered on the camera is projected into a bounded region of the image), and it must also be continuous at the camera center.

If we temporarily set aside the camera center, which we identify with the origin of  $\mathbb{R}^3$ , we may observe that a continuous extension over  $(\mathbb{R}^3)^*$  would have another problem: there would necessarily be points  $M$  that have zero reprojection error without actually being on the line of sight of point  $m$ .

To see this we consider the unit sphere  $\mathcal{S}$  centered on the origin. Let  $\mathcal{B}$  be the curve of camera-visible points in  $\mathcal{S}$  whose projection is the rectangular boundary of the image in  $\mathbb{R}^2$  (Fig. 1). This curve separates  $\mathcal{S}$  into two parts, one which is visible, and one which is hidden. Let us call the latter  $\mathcal{H}$ . On the one hand,  $\mathcal{H}$  is homeomorphic to a disk, thus simply connected, so its boundary  $\mathcal{B}$  can be contracted within  $\mathcal{H}$  to a point of  $\mathcal{B}$ ; and since we look for a projection function  $P$  that is continuous over  $(\mathbb{R}^3)^*$ , including over  $\mathcal{H}$ , the contraction path of  $\mathcal{B}$  within  $\mathcal{H}$  to a point of  $\mathcal{B}$  implies a contraction path of  $P(\mathcal{B})$  within  $P(\mathcal{H})$  to a point of  $P(\mathcal{B})$ . On the other hand,  $\mathbb{R}^2 \setminus \{m\}$  is not simply connected and the boundary of the image  $P(\mathcal{B})$  surrounds point  $m$  in  $\mathbb{R}^2$ , so it cannot be contracted to a point within  $\mathbb{R}^2 \setminus \{m\}$ , and *a fortiori* not within  $P(\mathcal{H}) \setminus \{m\}$ . Since  $P(\mathcal{B})$  can be contracted within  $P(\mathcal{H})$  and not within  $P(\mathcal{H}) \setminus \{m\}$ , these sets must differ, which implies  $m \in P(\mathcal{H})$ . Therefore there must exist at least one point in  $\mathcal{H}$  (thus invisible) whose projection is  $m$  (thus with zero reprojection error).

This observation can be extended to all spheres centered on the origin, so that we would have a family of non-visible points with zero reprojection error. To avoid unnecessary generalization we consider that this set of points would be a straight semi-line open at the origin and extending in some direction behind the camera.

### 2.3 Incidence function as a generalization of the reprojection error

The key idea is that since the projection function cannot be extended to the whole space in a continuous manner, nor with the property that the reprojection error is zero *only* for points  $M$  that are visible at  $m$ , we look for a more general function

$$F : \mathbb{R}^2 \times \mathbb{R}^3 \rightarrow \mathbb{R}^d \quad (1)$$

$$(m, M) \mapsto r$$

whose squared norm will be our cost function. This function assigns a residual vector  $r$  to a pair  $(m, M)$ . The residual vector,

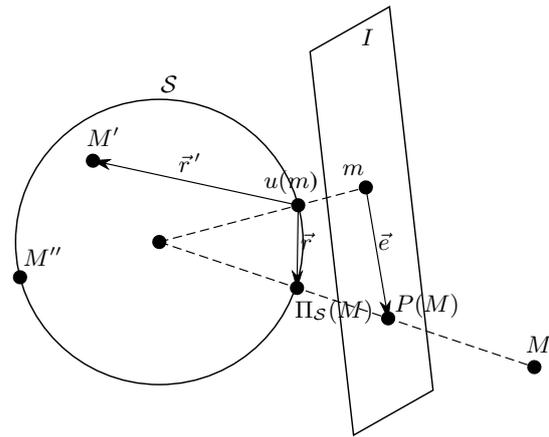


Figure 2. Sphere-based incidence function  $\Pi_{\mathcal{S}}(M) - u(m)$  between an image point  $m$  and an object point that is either outside ( $M$ ) or inside ( $M'$ ) the unit sphere. Case of a non-visible stationary point ( $M''$ ). Comparison of the 3D incidence residual vectors  $\vec{r}$  and  $\vec{r}'$  with the 2D reprojection error  $\vec{e} = P(M) - m$ .

whose dimension  $d$  is to be chosen, should be zero if and only if point  $M$  is on the line of sight of pixel  $m$ . For this reason we call  $F$  an incidence function. Given  $m$  we also define

$$F_m : M \mapsto F(m, M). \quad (2)$$

We want the incidence function  $F$  to have the following properties:

- It is defined for all  $M \in \mathbb{R}^3$  and  $m \in I$
- It is continuous and differentiable
- It is zero if and only if  $M$  is on the line of sight of  $m$
- $F_m$  is constant when  $M$  runs along a given line of sight
- The norm of  $F_m$  has no stationary point except when it is zero
- Near incidence (i.e.,  $F(m, M) \rightarrow 0$ ) it is equivalent to the reprojection error
- The norm of  $F_m$  is quasiconvex (optional)

### 2.4 Getting continuity at the camera center

It is obviously not possible for a non-constant function of  $M$  to be constant on each line of sight originating from the camera center and at the same time be continuous at the camera center. To have continuity we must relax the condition that the function should be constant on each line of sight. We start with the formulation

$$F(m, M) = \Pi_{\mathcal{S}}(M) - u(m) \quad (3)$$

where  $u(m) \in \mathcal{S}$  is the point of the unit sphere such that  $P(u) = m$ , and  $\Pi_{\mathcal{S}} : M \mapsto M / \max(1, |M|)$  centrally projects points on the unit sphere if they are outside, and leaves them unchanged otherwise (Fig. 2).

The function  $F_m$  has values in  $\mathbb{R}^3$ , is constant on each line of sight outside the unit sphere, and is zero if and only if point  $M$  is on the line of sight of  $m$  and outside the unit sphere.

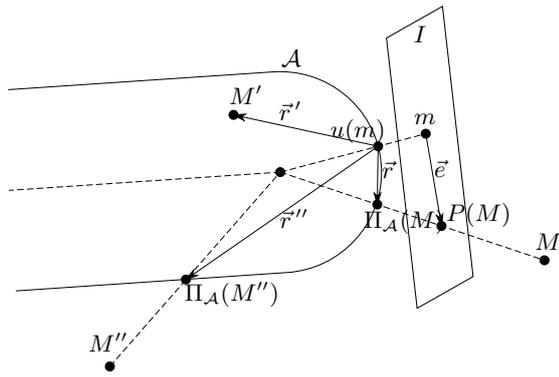


Figure 3. Half-cylinder-based incidence function between an image point  $m$  and object points in front of the camera ( $M$ ), inside surface  $\mathcal{A}$  ( $M'$ ), or behind the camera ( $M''$ ). Comparison of the 3D incidence residual vectors  $\vec{r}$ ,  $\vec{r}'$  and  $\vec{r}''$  with the 2D reprojection error  $\vec{e}$ .

There are still some issues: the norm of  $F_m$  has some non-zero stationary points, it is not equivalent to the norm of the reprojection error near incidence, and it is not quasiconvex. The first point can be seen by considering that outside of the unit sphere, the function can be expressed in terms of a function defined on the unit sphere, which is compact, so its norm must have a non-zero maximum which is a non-zero stationary point. The second point can be solved by applying a well-chosen linear transform to the residual vector, which we detail in section 2.6. The third point can be seen by observing that the set  $\{M : |F_m(M)| < \epsilon\}$  is not convex (also seen in later Fig. 5(a)).

## 2.5 Removing stationary points

It is desirable for the function  $F_m$  to not have any non-zero stationary point, in which case it can be safely minimized with any gradient-based technique. Also we would expect this to facilitate the optimization of a problem that is the sum of many such functions. It is not possible to avoid non-zero stationary points in the cost function if this function can be expressed in terms of the central projection of  $M$  on the unit sphere or any other compact surface. Conceptually, what we need to do is remove one point from the unit sphere, or the associated oriented direction in space, where the cost function would be stationary.

So we define an open surface  $\mathcal{A}$  that has one intersection with each semi-line originating from the camera center except for the removed direction. We define  $\Pi_{\mathcal{A}}(M)$  as the central projection of point  $M$  on surface  $\mathcal{A}$  when  $M$  is farther to the origin than the surface, or  $M$  otherwise. We also update the previous definition of  $u(m)$  such that it now represents the point on  $\mathcal{A}$  (instead of  $S$ ) whose image by  $P$  is  $m$ . Then we define

$$F(m, M) = \Pi_{\mathcal{A}}(M) - u(m). \quad (4)$$

For this formulation to be continuous, surface  $\mathcal{A}$  should stretch towards infinity when it tends towards the singular direction (it cannot simply be a punctured sphere). There are several possibilities, each with their advantages and disadvantages in terms of mathematical properties or computational efficiency. One solution is to define  $\mathcal{A}$  as the union of a semi-sphere in front of the camera and a semi-cylinder behind the camera (Fig. 3). It has the advantage of being quite simple. It also illustratively isolates the singularity at the camera center and the singularity in the direction space. One drawback, that is shared by all

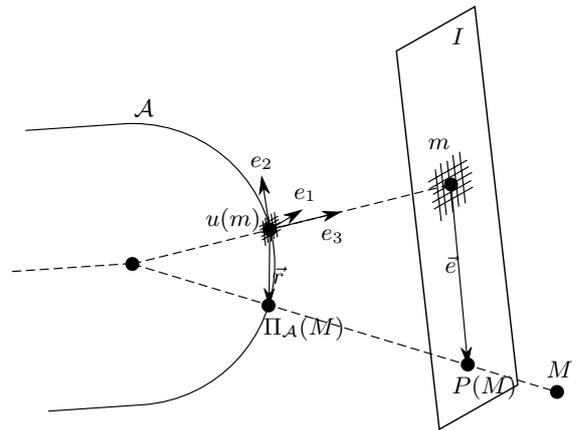


Figure 4. Using an appropriate vector basis  $(e_1, e_2, e_3)$  on  $\mathbb{R}^3$  to measure the 3D incidence residual vector  $\vec{r}$  makes it equivalent to the pixel reprojection error  $\vec{e}$ .

choices of  $\mathcal{A}$ , is that a radius for the sphere and cylinder has to be chosen more or less arbitrarily.

Another choice would be a paraboloid, and if we extend the set of directions forbidden for projection, we could also use one sheet of an hyperboloid, or even a plane at some distance in front of the camera. Finally we could define surface  $\mathcal{A}$  to be dependent on  $m$  (for instance choosing the forbidden direction opposite to the line of sight of  $m$ , instead of the optical axis of the camera).

## 2.6 Equivalence to the reprojection error

In order to use the incidence function in standard bundle adjustment interchangeably with the reprojection error, we need it to be equivalent to the reprojection error when it tends towards zero. For this all we need is that for any given  $m$  the Jacobians of  $M \mapsto P(M) - m$  and  $M \mapsto F_m(M)$  be equal when  $F_m(M) = 0$ . We first note that the reprojection error has values in  $\mathbb{R}^2$  whereas the incidence function has values in  $\mathbb{R}^3$ . However, outside of  $\mathcal{A}$ ,  $F_m(M)$  has really values in  $\mathcal{A} - u(m)$ , which is a 2-manifold, so it should be possible to find a local linear mapping between  $F_m(M)$  and  $P(M) - m$  augmented with a zero third component (Fig. 4).

Thus, we look for an invertible  $3 \times 3$  matrix  $K$  such that under the condition  $F_m(M) = 0$  we have

$$\begin{bmatrix} \frac{\partial P}{\partial M} \\ 0 \end{bmatrix} = K \frac{\partial F_m}{\partial M}. \quad (5)$$

Matrices  $\frac{\partial P}{\partial M}$  and  $\frac{\partial F_m}{\partial M}$  have dimensions  $2 \times 3$  and  $3 \times 3$  respectively. They both have the same one-dimensional right null space, defined by the direction of the line of sight of  $m$ . They also both have rank 2. Given that  $\frac{\partial P}{\partial M}$  has full rank, contrary to  $\frac{\partial F_m}{\partial M}$ , establishing the solution for  $L = K^{-1}$  instead of  $K$ , happens to be more straightforward. If we note  $L_{12}$  the  $L$  submatrix composed of its first two columns, and  $L_3$  its last column, then multiplying (5) by  $L$  on the left we obtain

$$[L_{12}L_3] \begin{bmatrix} \frac{\partial P}{\partial M} \\ 0 \end{bmatrix} = L_{12} \frac{\partial P}{\partial M} = \frac{\partial F_m}{\partial M}. \quad (6)$$

Multiplying (6) on the right by the pseudo inverse of full rank

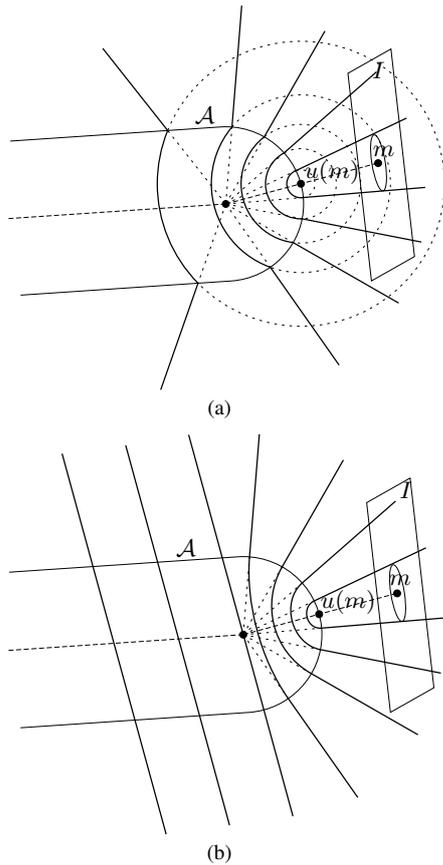


Figure 5. (a) Iso-level surfaces of  $|F_m|$  (bold solid lines). (b) Design of iso-level surfaces for a quasiconvex function.

matrix  $\frac{\partial P}{\partial M}$  we get

$$L_{12} = \frac{\partial F_m}{\partial M} \left( \frac{\partial P}{\partial M} \right)^T \left( \frac{\partial P}{\partial M} \left( \frac{\partial P}{\partial M} \right)^T \right)^{-1}. \quad (7)$$

Note that for practical computations, any vector basis of  $\mathbb{R}^3$  for the derivation variable  $M$  can be chosen, for instance one in which the third axis is aligned with the line of sight of  $m$ . In this case the third column of  $\frac{\partial F_m}{\partial M}$  and  $\frac{\partial P}{\partial M}$  is zero. Then the formula can be simplified by deriving only with respect to any other two independent axes, and using a plain matrix inverse instead of a pseudo-inverse.

As  $L_3$  has vanished from the equations, it can be chosen arbitrarily, but we need  $L$  (and  $K$ ) to be full rank to capture the third dimension of the residual vector whenever we are further from incidence. So we set  $L_3$  as (a multiple of) the cross product of the column vectors of  $L_{12}$ , and then compute  $K = L^{-1}$ .

Finally the modified incidence function is

$$G(m, M) = K(m)F(m, M) \quad (8)$$

where  $K(m)$  is the matrix defined above which depends only on  $m$  (it can be evaluated at any chosen point  $M$  such that  $F_m(M) = 0$ , for instance at  $M = u(m)$ ).

### 2.7 Towards a quasiconvex cost function

The cost function defined so far is not quasiconvex. It could be made quasiconvex by designing level surfaces that enclose

convex volumes.

As pointed out in (Kahl and Hartley, 2008) it would be interesting to have quasiconvex cost functions if we would optimize their point-wise maximum, as quasiconvexity is preserved by the point-wise maximum operator. However, in the context of standard bundle adjustment (using not the maximum, but the sum of individual cost functions) quasiconvexity brings no useful property, so we did not go so far as to actually make our incidence function quasiconvex. Convexity would have been even more desirable, as it is preserved by addition, however the squared norm of the reprojection error is not itself convex (Kahl and Hartley, 2008) so the squared norm of a function equivalent to the reprojection error would be non-convex either.

This being said, we provide an example of what a quasiconvex cost function would look like, by modifying the sphere-and-cylinder-based one (Fig. 5). A way to adjust an existing incidence function to fit designed iso-level surfaces, is to first determine which iso-level surface of  $F_m$  a given point  $M$  is on, with the associated value, and then scale the residual vector so that its length has this value.

### 2.8 Remarks

An important property of the incidence function is that contrary to the reprojection error, it does not depend on the definition, or extrapolation, of the image geometrical model away from point  $m$ . This is especially useful in an image featuring high radial distortion modeled by a polynomial function, because in this case evaluating the projection function for a point  $M$  that is far from being visible at point  $m$  may fail or provide a hazardous result.

It may also be useful to apply the incidence function idea in the context of more complex geometrical models, such as images acquired by a pushbroom sensor. In this case, not only is the projection function costly to evaluate, but the geometry may present many accidents (especially in aerial pushbroom) such as a same object being imaged more than once, or the image being jittery, perturbing the evaluation of the reprojection error and of its gradient. When evaluating the reprojection error between  $M$  and  $m$ , we do not want to be perturbed by any imaging accidents located far from  $m$  and otherwise irrelevant. Like for the pinhole model, we can achieve this by defining an incidence function between  $M$  and  $m$  that depends only on differential properties of the projection at  $m$ .

## 3. EXPERIMENTAL RESULTS

### 3.1 Point triangulation

Using the new cost function, we verified that it is possible to use standard optimization algorithms to triangulate the position of a point seen in multiple images starting from an arbitrary initial position. There is the risk of converging to a local minimum, but this is no different than when using the quadratic sum of reprojection errors which features such local minima (Kahl and Hartley, 2008). When points are behind cameras the value of the incidence function is large, so that they are naturally drawn in front of the cameras. However the squared sum of residuals could conceivably have local minima there too depending on the configuration of the cameras.

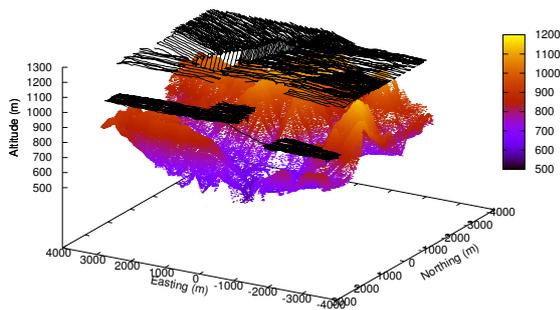


Figure 6. Aerial acquisition over mountainous area. Recovered trajectory (black line) and tie point positions (colored dots).

### 3.2 Bundle adjustment

Seeing that the cost function allows to correctly triangulate points, as if the reprojection error were minimized, the next step is to plug it into a real bundle adjustment problem and check that it also gets the correct solution.

As with all bundle adjustment, the assumption is that the initial camera orientation is close enough to the optimum. We initialize points using this geometry and the triangulation algorithm described previously (section 3.1). Then we optimize the full bundle adjustment problem, solving both for camera parameters (exterior and interior orientation) and point positions using the Levenberg-Marquard algorithm with sparse Cholesky factorization of the reduced system Gauss-Newton Hessian approximation.

We verified on several examples that, as expected, the obtained solution is the same whether we optimize the reprojection error or the incidence function. Outlier tie points may have significantly different costs between the two algorithms, thus leading to slightly different solutions. In this case the one obtained with the incidence function is likely less perturbed by gross outliers as this function is bounded, but in the end these outliers are filtered and final results are the same.

### 3.3 Relative pose with known rotation

Estimating the relative pose between cameras is a fundamental problem for which several algorithms have been proposed. One class of methods consists in estimating first the relative orientations, by computing pairwise orientations and integrating them (Chatterjee and Govindu, 2013), and then estimating the relative positions with fixed orientations (Moulon et al., 2013).

We propose to use the bundle adjustment engine based on the new cost function to solve the known-rotation relative position problem. For this we initialize the position of all cameras at a single arbitrary location and optimize the global cost function. The properties of the new cost function may enable the process to converge to a realistic solution despite the initial geometry being far from the optimum. This is what we empirically observe in a variety of configurations.

We present results obtained with two different configurations: a bidimensional aerial acquisition, and a mono-dimensional multi-camera mobile-mapping-system acquisition.

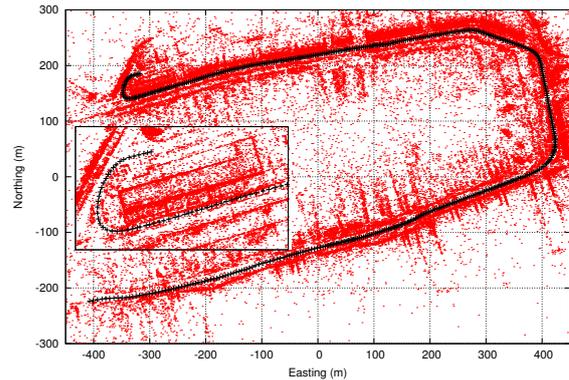


Figure 7. Ground acquisition in Marseille port. Recovered trajectory (black line) and tie point positions (red dots), with a closeup of one extremity.

**3.3.1 Aerial acquisition** This dataset consists of 22,123 images taken from a Unmanned Aerial Vehicle (UAV) over a  $7\text{ km} \times 7\text{ km}$  mountainous area following a structured flight plan. The images possess GPS position and IMU orientation information. They also have been finely bundle adjusted to obtain camera calibration, known relative orientations, and some kind of “ground-truth” against which to compare the result of our relative position algorithm.

A set of about 216,000 tie points has been obtained by image block matching in cartographic space, with 146 points per image on average. Any outlier has been identified and removed during the bundle adjustment process so that the relative position experiment may be conducted in ideal conditions.

We initialized image positions and tie point positions to be at the origin of the coordinate system. As for camera calibration and relative orientations we fixed them to the values obtained previously by bundle adjustment.

Then we optimized the incidence-based cost function using the Levenberg-Marquard algorithm, which yielded a solution for relative image positions and tie point positions. We registered the computed relative positions to the ground-truth positions, using the similarity transform that minimizes the quadratic sum of point-wise distances. Then we measured the distance (or error) between the ground-truth positions and the registered relative positions. The median error distance is 0.68 m and 97% are under 2 m. The result can be seen in Fig. 6.

So, for these images the process converged well, and to a sensible solution. Yet we may wonder why it did not get closer to the bundle adjustment result, especially as we used the camera calibration and relative orientations given by this very bundle adjustment, with only the relative positions left to be estimated. A factor to consider is that the original bundle adjustment problem has a prior cost term based on given GPS positions that makes it substantially different than the relative position problem where no such term exists.

**3.3.2 Ground acquisition** This dataset consists of 526 shots taken by a multi-camera system mounted on a car. The acquisition system is made of 6 cameras mounted at  $60^\circ$  from each other with a little overlap, effectively covering a  $360^\circ$  cylindrical field of view. The maximum distance between any two camera optical centers is about 50 cm. The car ran along the

2 km-long street bordering the port of Marseille, France. Each shot came with unreliable GPS position and heading angle.

For bundle adjustment, a collection of about 1.5 million SIFT-like (Lowe, 2004) tie points has been obtained between all 3,156 images. The result serves as “ground-truth”, and provides geometrical calibration of the system, as well as an estimate of the relative orientation between shots.

We used the same protocol as for the aerial dataset, and initialized the positions of tie points to the center of the coordinate system, as well as the positions of the camera rig for the 526 shots. Then we optimized the incidence-based cost function, and registered the obtained rig positions to the “ground-truth” with the best similarity transform. Fig. 7 shows the obtained results.

The measured median distance error is 0.61 m with 90% under 2 m. This implies that there was relatively little drift in spite of the long linear acquisition. The known rotations certainly helped with that, but the configuration probably did too. In fact the inside of the port being relatively flat and unobstructed, both sides of the port could virtually be seen from anywhere, providing some correspondances and rigidity between all images.

Worth of note is that the scale factor of the similarity used to register the solution to the ground-truth is 1.001, which means that the unregistered solution was already recovered with almost the correct scale. This is only possible because we are not dealing with a pinhole camera system but with a generalized camera system, whose calibrated inter-camera distances provide an absolute scale reference.

**3.3.3 Discussion** The results presented so far show that the incidence-based cost function provide robustness within a bundle adjustment algorithm. Although not the primary goal, we found that computing unknown relative positions was even possible using standard bundle adjustment with the new cost function.

It should be noted though that for this experiment we placed ourselves in ideal conditions: calibrated camera system, fairly accurate relative orientation taken from bundle adjustment, and outlier-free tie points.

When these conditions are not met, the process would still converge, but would not provide the right positions. It is still possible to use the proposed bundle adjustment engine as a building block to iteratively refine the internal camera geometry, or image relative orientation, or filter outliers, but that is outside the scope of this study.

Even in ideal conditions, there are some cases in which the process does not converge to a sensible solution. It is especially so when the connectivity graph of the images induced by the tie points has several connected components, or when there are some very weak links, or when the links are very unbalanced. In such cases we can end up with a solution that is globally wrong, but that is piece-wise correct (only the relative position and scale between those pieces are wrong). Linear acquisitions (along a street or air corridor) are more prone to this risk than area acquisitions.

Also some numerical difficulties may arise due to a kind of tie point that would be seen by almost collinear rays in different images (even if this is a transitory accident during optimization). Such points may get triangulated very far away and end

up being stuck there when the Jacobian of their cost function becomes virtually zero. This may require specific handling in the optimization algorithm, or taking it into account in the design of the incidence function.

Finally, the value of the radius used to isolate the singularity in the incidence function can have an influence, especially when the problem has an absolute scale factor, such as with standard bundle adjustment, or with relative positions of a generalized camera system. In these cases the radius must absolutely be smaller than the object-to-camera distances. Other than that, it does not influence the optimum but it can still influence the optimization path.

#### 4. CONCLUSION

We presented an alternative formulation of the bundle adjustment problem that is equivalent to minimizing the quadratic sum of reprojection errors but that is defined for all point positions and is exempt of singularities. It allows an easier implementation without special cases and enjoys a larger radius of convergence, even so as to provide a new solution to the known-orientation relative position estimation problem.

Further work may be geared towards improving the mathematical properties of the incidence function to allow tackling an even wider range of problems. For instance it could be made quasiconvex and used with  $L_\infty$ -norm thus yielding provable convergence for some optimization problems. The incidence function idea could also be applied to other sensors such as optical pushbroom. Finally, some comparative experiments could be carried out on benchmark datasets.

#### REFERENCES

- Chatterjee, A., Govindu, V. M., 2013. Efficient and Robust Large-Scale Rotation Averaging. *2013 IEEE International Conference on Computer Vision*, 521-528.
- Kahl, F., Hartley, R. I., 2008. Multiple View Geometry Under the  $L_\infty$ -norm. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Lowe, D. G., 2004. Distinctive Image Features from Scale-Invariant Keypoints. *Int. J. Comput. Vision*, 60(2), 91–110. <https://doi.org/10.1023/B:VISI.0000029664.99615.94>.
- Mitra, K., Chellappa, R., 2008. A scalable projective bundle adjustment algorithm using the  $l_\infty$  norm. *Sixth Indian Conference on Computer Vision, Graphics & Image Processing (ICVGIP) 2008, Bhubaneswar, India, 16-19 December 2008*, 79–86.
- Moulon, P., Monasse, P., Marlet, R., 2013. Global fusion of relative motion for robust, accurate and scalable structure from motion. *ICCV*, IEEE Computer Society, 3248–3255.
- Schinstock, D., Lewis, C., Buckley, C., 2009. An alternative cost function to bundle adjustment used for aerial photography from UAVs. *ASPRS Annual Conference*.
- Triggs, B., McLauchlan, P. F., Hartley, R. I., Fitzgibbon, A. W., 2000. Bundle adjustment - a modern synthesis. *Proceedings of the International Workshop on Vision Algorithms: Theory and Practice*, ICCV '99, Springer-Verlag, London, UK, UK, 298–372.